

Translating Answer-Set Programs into Bit-Vector Logic^{*}

Mai Nguyen, Tomi Janhunen, and Ilkka Niemelä

Aalto University School of Science
Department of Information and Computer Science
{Mai.Nguyen,Tomi.Janhunen,Ilkka.Niemela}@aalto.fi

Abstract. Answer set programming (ASP) is a paradigm for declarative problem solving where problems are first formalized as rule sets, i.e., answer-set programs, in a uniform way and then solved by computing answer sets for programs. The satisfiability modulo theories (SMT) framework follows a similar modelling philosophy but the syntax is based on extensions of propositional logic rather than rules. Quite recently, a translation from answer-set programs into difference logic was provided—enabling the use of particular SMT solvers for the computation of answer sets. In this paper, the translation is revised for another SMT fragment, namely that based on fixed-width bit-vector theories. Thus, even further SMT solvers can be harnessed for the task of computing answer sets. The results of a preliminary experimental comparison are also reported. They suggest a level of performance which is similar to that achieved via difference logic.

1 Introduction

Answer set programming (ASP) is a rule-based approach to declarative problem solving [15, 22, 24]. The idea is to first formalize a given problem as a set of rules also called an *answer-set program* so that the answer sets of the program correspond to the solution of the problem. Such problem descriptions are typically devised in a *uniform* way which distinguishes general principles and constraints of the problem in question from any instance-specific data. To this end, term variables are deployed for the sake of compact representation of rules. Solutions themselves can then be found out by *grounding* the rules of the answer-set program, and by computing answer sets for the resulting ground program using an answer set solver. State-of-the-art answer set solvers are already very efficient search engines [7, 11] and have a wide range of industrial applications.

The satisfiability modulo theories (SMT) framework [3] follows a similar modelling philosophy but the syntax is based on extensions of propositional logic rather than rules with term variables. The SMT framework enriches traditional satisfiability (SAT) checking [5] in terms of background theories which are selected amongst a number of alternatives.¹ Parallel to propositional atoms, also *theory atoms* involving non-Boolean variables² can be used as references to potentially infinite domains. Theory atoms are typically used to express various constraints such as linear constraints, difference constraints, etc., and they enable very concise representations of certain problem domains for which plain Boolean logic would be more verbose or insufficient in the first place.

As regards the relationship of ASP and SMT, it was quite recently shown [20, 25] that answer-set programs can be efficiently translated into a simple SMT fragment, namely *difference logic* (DL) [26]. This fragment is based on theory atoms of the form $x - y \leq k$ formalizing an upper bound k on the *difference* of two integer-domain variables x and y . Although the required transformation is linear, it is not reasonable to expect that such theories are directly written by humans in order to express the essentials of ASP in SMT. The translations from [20, 25] and their implementation called LP2DIFF³ enable the use of particular SMT solvers for the computation of answer sets. Our experimental results [20] indicate that the performance obtained in this way is surprisingly close to that of state-of-the-art answer set solvers. The

^{*} This paper appears in the Proceedings of the 19th International Conference on Applications of Declarative Programming and Knowledge Management (INAP 2011).

¹ <http://combination.cs.uiowa.edu/smtlib/>

² However, variables in SMT are syntactically represented by (functional) constants having a free interpretation over a specific domain such as integers or reals.

³ <http://www.tcs.hut.fi/Software/lp2diff/>

results of the third ASP competition [7], however, suggest that the performance gap has grown since the previous competition. To address this trend, our current and future agendas include a number of points:

- We gradually increase the number of supported SMT fragments which enables the use of further SMT solvers for the task of computing answer sets.
- We continue the development of new translation techniques from ASP to SMT.
- We submit ASP-based benchmark sets to future SMT competitions (SMT-COMPs) to foster the efficiency of SMT solvers on problems that are relevant for ASP.
- We develop new integrated languages that combine features of ASP and SMT, and aim at implementations via translation into pure SMT as initiated in [18].

This paper contributes to the first item by devising a translation from answer-set programs into theories of bit-vector logic. There is a great interest to develop efficient solvers for this particular SMT fragment due to its industrial relevance. In view of the second item, we generalize an existing translation from [20] to the case of bit-vector logic. Using an implementation of the new translation, viz. LP2BV, new benchmark classes can be created to support the third item on our agenda. Finally, the translation also creates new potential for language integration. In the long run, rule-based languages and, in particular, the modern grounders exploited in ASP can provide valuable machinery for the generation of SMT theories in analogy to answer-set programs: The *source code* of an SMT theory can be compacted using rules and term variables [18] and specified in a uniform way which is independent of any concrete problem instances. Analogous approaches [2, 14, 23] combine ASP and constraint programming techniques without a translation.

The rest of this paper is organized as follows. First, the basic definitions and concepts of answer-set programs and fixed-width bit-vector logic are briefly reviewed in Section 2. The new translation from answer-set programs into bit-vector theories is then devised in Section 3. The extended rule types of SMOELS compatible systems are addressed in Section 4. Such extensions can be covered either by native translations into bit-vector logic or translations into normal programs. As part of this research, we carried out a number of experiments using benchmarks from the second ASP competition [11] and two state-of-the-art SMT solvers, viz. BOOLECTOR and Z3. The results of the experiments are reported in Section 5. Finally, we conclude this paper in Section 6 in terms of discussions of results and future work.

2 Preliminaries

The goal of this section is to briefly review the source and target formalisms for the new translation devised in the sequel. First, in Section 2.1, we recall normal logic programs subject to answer set semantics and the main notions exploited in their translation. A formal account of bit-vector logic follows in Section 2.2.

2.1 Normal Logic Programs

As usual, we define a *normal logic program* P as a finite set of *rules* of the form

$$a \leftarrow b_1, \dots, b_n, \sim c_1, \dots, \sim c_m \quad (1)$$

where a, b_1, \dots, b_n , and c_1, \dots, c_m are propositional atoms and \sim denotes *default negation*. The *head* of a rule r of the form (1) is $\text{hd}(r) = a$ whereas the part after the symbol \leftarrow forms the *body* of r , denoted by $\text{bd}(r)$. The body $\text{bd}(r)$ consists of the positive part $\text{bd}^+(r) = \{b_1, \dots, b_n\}$ and the negative part $\text{bd}^-(r) = \{c_1, \dots, c_m\}$ so that $\text{bd}(r) = \text{bd}^+(r) \cup \{\sim c \mid c \in \text{bd}^-(r)\}$. Intuitively, a rule r of the form (1) appearing in a program P is used as follows: the head $\text{hd}(r)$ can be inferred by r if the *positive body atoms* in $\text{bd}^+(r)$ are inferable by the other rules of P , but not the *negative body atoms* in $\text{bd}^-(r)$. The positive part of the rule, r^+ is defined as $\text{hd}(r) \leftarrow \text{bd}^+(r)$. A normal logic program is called *positive* if $r = r^+$ holds for every rule $r \in P$.

Semantics To define the semantics of a normal program P , we let $\text{At}(P)$ stand for the set of atoms that appear in P . An *interpretation* of P is any subset $I \subseteq \text{At}(P)$ such that for an atom $a \in \text{At}(P)$, a is *true* in I , denoted $I \models a$, iff $a \in I$. For any negative literal $\sim c$, $I \models \sim c$ iff $I \not\models c$ iff $c \notin I$. A rule r is satisfied in I , denoted $I \models r$, iff $I \models \text{bd}(r)$ implies $I \models \text{hd}(r)$. An interpretation I is a *classical model* of P , denoted $I \models P$, iff, $I \models r$ holds for every $r \in P$. A model $M \models P$ is a *minimal model* of P iff there is no $M' \models P$ such that $M' \subset M$. Each positive normal program P has a unique minimal model, i.e., the *least model* of P denoted by $\text{LM}(P)$ in the sequel. The least model semantics can be extended for an arbitrary normal program P by *reducing* P into a positive program $P^M = \{r^+ \mid r \in P \text{ and } M \cap \text{bd}^-(r) = \emptyset\}$ with respect to $M \subseteq \text{At}(P)$. Then *answer sets*, also known as *stable models* [16], can be defined.

Definition 1 (Gelfond and Lifschitz [16]). An interpretation $M \subseteq \text{At}(P)$ is an answer set of a normal program P iff $M = \text{LM}(P^M)$.

Example 1. Consider a normal program P [20] consisting of the following six rules:

$$\begin{array}{lll} a \leftarrow b, c. & a \leftarrow d. & b \leftarrow a, \sim d. \\ b \leftarrow a, \sim c. & c \leftarrow \sim d. & d \leftarrow \sim c. \end{array}$$

The answer sets of P are $M_1 = \{a, b, d\}$ and $M_2 = \{c\}$. To verify the latter, we note that $P^{M_2} = \{a \leftarrow b, c; b \leftarrow a; c \leftarrow \sim d; a \leftarrow d\}$ for which $\text{LM}(P^{M_2}) = \{c\}$. On the other hand, we have $P^{M_3} = P^{M_2}$ for $M_3 = \{a, b, c\}$ so that $M_3 \notin \text{AS}(P)$. ■

The number of answer sets possessed by a normal program P can vary in general. The set of answer sets of a normal program P is denoted by $\text{AS}(P)$. Next we present some concepts and results that are relevant in order to capture answer sets in terms of propositional logic and its extensions in the SMT framework.

Completion Given a normal program P and an atom $a \in \text{At}(P)$, the *definition* of a in P is the set of rules $\text{Def}_P(a) = \{r \in P \mid \text{hd}(r) = a\}$. The *completion* of a normal program P , denoted by $\text{Comp}(P)$, is a propositional theory [8] which contains

$$a \leftrightarrow \bigvee_{r \in \text{Def}_P(a)} \left(\bigwedge_{b \in \text{bd}^+(r)} b \wedge \bigwedge_{c \in \text{bd}^-(r)} \neg c \right) \quad (2)$$

for each atom $a \in \text{At}(P)$. Given a propositional theory T and its signature $\text{At}(T)$, the semantics of T is determined by $\text{CM}(T) = \{M \subseteq \text{At}(T) \mid M \models T\}$. It is possible to relate $\text{CM}(\text{Comp}(P))$ with the models of a normal program P by distinguishing *supported models* [1] for P . A model $M \models P$ is a supported model of P iff for every atom $a \in M$ there is a rule $r \in P$ such that $\text{hd}(r) = a$ and $M \models \text{bd}(r)$. In general, the set of supported models $\text{SuppM}(P)$ of a normal program P coincides with $\text{CM}(\text{Comp}(P))$. It can be shown [21] that stable models are also supported models but not necessarily vice versa. This means that in order to capture $\text{AS}(P)$ using $\text{Comp}(P)$, the latter has to be extended in terms of additional constraints as done, e.g., in [17, 20].

Example 2. For the program P of Example 1, the theory $\text{Comp}(P)$ has formulas $a \leftrightarrow (b \wedge c) \vee d$, $b \leftrightarrow (a \wedge \neg d) \vee (a \wedge \neg c)$, $c \leftrightarrow \neg d$, and $d \leftrightarrow \neg c$. The models of $\text{Comp}(P)$, i.e., its supported models, are $M_1 = \{a, b, d\}$, $M_2 = \{c\}$, and $M_3 = \{a, b, c\}$. ■

Dependency Graphs The *positive dependency graph* of a normal program P , denoted by $\text{DG}^+(P)$, is a pair $\langle \text{At}(P), \leq \rangle$ where $b \leq a$ holds iff there is a rule $r \in P$ such that $\text{hd}(r) = a$ and $b \in \text{bd}^+(r)$. Let \leq^* denote the *reflexive and transitive closure* of \leq . A *strongly connected component* (SCC) of $\text{DG}^+(P)$ is a maximal non-empty subset $S \subseteq \text{At}(P)$ such that $a \leq^* b$ and $b \leq^* a$ hold for each $a, b \in S$. The set of defining rules is generalized for an SCC S by $\text{Def}_P(S) = \bigcup_{a \in S} \text{Def}_P(a)$. This set can be naturally partitioned into sets $\text{Ext}_P(S) = \{r \in \text{Def}_P(S) \mid \text{bd}^+(r) \cap S = \emptyset\}$ and $\text{Int}_P(S) = \{r \in \text{Def}_P(S) \mid \text{bd}^+(r) \cap S \neq \emptyset\}$ of *external* and *internal* rules associated with S , respectively. Thus, $\text{Def}_P(S) = \text{Ext}_P(S) \sqcup \text{Int}_P(S)$ holds in general.

Example 3. In the case of the program P from Example 1, the SCCs of $\text{DG}^+(P)$ are $S_1 = \{a, b\}$, $S_2 = \{c\}$, and $S_3 = \{d\}$. For S_1 , we have $\text{Ext}_P(S_1) = \{a \leftarrow d\}$. ■

2.2 Bit-Vector Logic

Fixed-width bit-vector theories have been introduced for high-level reasoning about digital circuitry and computer programs in the SMT framework [27, 4]. Such theories are expressed in an extension of propositional logic where atomic formulas speak about bit vectors in terms of a rich variety of operators.

Syntax As usual in the context of SMT, variables are realized as constants that have a free interpretation over a particular domain (such as integers or reals)⁴. In the case of fixed-width bit-vector theories, this means that each constant symbol x represents a vector $x[1 \dots m]$ of bits of particular width m , denoted by $w(x)$ in the sequel. Such vectors enable a more compact representation of structures like registers and often allow more efficient reasoning about them. A special notation \bar{n} is introduced to denote a bit vector that equals to n , i.e., \bar{n} provides a binary representation of n . We assume that the actual width $m \geq \log_2(n+1)$ is determined by the context where the notation \bar{n} is used. For the purposes of this paper, the most interesting arithmetic operator for combining bit vectors is the addition of two m -bit vectors, denoted by the parameterized function symbol $+_m$ in an infix notation. The resulting vector is also m -bit which can lead to an overflow if the sum exceeds $2^m - 1$. Moreover, we use Boolean operators $=_m$ and $<_m$ with the usual meanings for comparing the values of two m -bit vectors. Thus, assuming that x and y are m -bit free constants, we may write atomic formulas like $x =_m y$ and $x <_m y$ in order to compare the m -bit values of x and y . In addition to syntactic elements mentioned so far, we can use the primitives of propositional logic to build more complex *well-formed formulas* of bit-vector logic. The syntax defined for the SMT library contains further primitives which are skipped in this paper. A theory T in bit-vector logic is a set of well-formed bit-vector formulas as illustrated by the following example.

Example 4. Consider a system of two processes, say A and B, and a theory $T = \{a \rightarrow (x <_2 y), b \rightarrow (y <_2 x)\}$ formalizing a scheduling policy for them. The intuitive reading of a (resp. b) is that process A (resp. B) is scheduled with a higher priority and, thus, should start earlier. The constants x and y denote the respective starting times of A and B. Thus, e.g., $x <_2 y$ means that process A starts before process B. ■

Semantics Given a bit-vector theory T , we write $\text{At}(T)$ and $\text{FC}(T)$ for the sets of propositional atoms and free constants, respectively, appearing in T . To determine the semantics of T , we define *interpretations* for T as pairs $\langle I, \tau \rangle$ where $I \subseteq \text{At}(T)$ is a standard propositional interpretation and τ is a partial function that maps a free constant $x \in \text{FC}(T)$ and an index $1 \leq i \leq w(x)$ to the set of bits $\{0, 1\}$. Given τ , a constant $x \in \text{FC}(T)$ is mapped onto $\tau(x) = \sum_{i=1}^{w(x)} (\tau(x, i) \cdot 2^{w(x)-i})$ and, in particular, $\tau(\bar{n}) = n$ for any n . To cover any *well-formed terms*⁵ t_1 and t_2 involving $+_m$ and m -bit constants from $\text{FC}(T)$, we define $\tau(t_1 +_m t_2) = \tau(t_1) + \tau(t_2) \bmod 2^m$ and $w(t_1 +_m t_2) = m$. Hence, the value $\tau(t)$ can be determined for any well-formed term t which enables the evaluation of more complex formulas as formalized below.

Definition 2. Let T be a bit-vector theory, $a \in \text{At}(T)$ a propositional atom, t_1 and t_2 well-formed terms over $\text{FC}(T)$ such that $w(t_1) = w(t_2)$, and ϕ and ψ well-formed formulas. Given an interpretation $\langle I, \tau \rangle$ for the theory T , we define

1. $\langle I, \tau \rangle \models a \iff a \in I$,
2. $\langle I, \tau \rangle \models t_1 =_m t_2 \iff \tau(t_1) = \tau(t_2)$,
3. $\langle I, \tau \rangle \models t_1 <_m t_2 \iff \tau(t_1) < \tau(t_2)$,
4. $\langle I, \tau \rangle \models \neg \phi \iff \langle I, \tau \rangle \not\models \phi$,
5. $\langle I, \tau \rangle \models \phi \vee \psi \iff \langle I, \tau \rangle \models \phi \text{ or } \langle I, \tau \rangle \models \psi$,
6. $\langle I, \tau \rangle \models \phi \rightarrow \psi \iff \langle I, \tau \rangle \not\models \phi \text{ or } \langle I, \tau \rangle \models \psi$, and
7. $\langle I, \tau \rangle \models \phi \leftrightarrow \psi \iff \langle I, \tau \rangle \models \phi \text{ if and only if } \langle I, \tau \rangle \models \psi$.

The interpretation $\langle I, \tau \rangle$ is a model of T , i.e., $\langle I, \tau \rangle \models T$, iff $\langle I, \tau \rangle \models \phi$ for all $\phi \in T$.

⁴ We use typically symbols x, y, z to denote such free (functional) constants and symbols a, b, c to denote propositional atoms.

⁵ The constants and operators appearing in a well-formed term t are based on a fixed width m . Moreover, the width $w(x)$ of each constant $x \in \text{FC}(T)$ must be the same throughout T .

It is clear by Definition 2 that pure propositional theories T are treated classically, i.e., $\langle I, \tau \rangle \models T$ iff $I \models T$ in the sense of propositional logic. As regards the theory T from Example 4, we have the sets of symbols $\text{At}(T) = \{a, b\}$ and $\text{FC}(T) = \{x, y\}$. Furthermore, we observe that there is no model of T of the form $\langle \{a, b\}, \tau \rangle$ because it is impossible to satisfy $x <_2 y$ and $y <_2 x$ simultaneously using any partial function τ . On the other hand, there are 6 models of the form $\langle \{a\}, \tau \rangle$ because $x <_2 y$ can be satisfied in $3 + 2 + 1 = 6$ ways by picking different values for the 2-bit vectors x and y .

3 Translation

In this section, we present a translation of a logic program P into a bit-vector theory $\text{BV}(P)$ that is similar to an existing translation [20] into difference logic. As its predecessor, the translation $\text{BV}(P)$ consists of two parts. Clark's completion [8], denoted by $\text{CC}(P)$, forms the first part of $\text{BV}(P)$. The second part, i.e., $\text{R}(P)$, is based on *ranking constraints* from [25] so that $\text{BV}(P) = \text{CC}(P) \cup \text{R}(P)$. Intuitively, the idea is that the completion $\text{CC}(P)$ captures *supported models* of P [1] and the further formulas in $\text{R}(P)$ exclude the non-stable ones so that any classical model of $\text{BV}(P)$ corresponds to a stable model of P .

The completion $\text{CC}(P)$ is formed for each atom $a \in \text{At}(P)$ on the basis of (2):

1. If $\text{Def}_P(a) = \emptyset$, the formula $\neg a$ is included to capture the corresponding empty disjunction in (2).
2. If there is $r \in \text{Def}_P(a)$ such that $\text{bd}(r) = \emptyset$, then one of the disjuncts in (2) is trivially true and the formula a can be used as such to capture the definition of a .
3. If $\text{Def}_P(a) = \{r\}$ for a rule $r \in P$ with $n + m > 0$, then we simplify (2) to a formula of the form

$$a \leftrightarrow \bigwedge_{b \in \text{bd}^+(r)} b \wedge \bigwedge_{c \in \text{bd}^-(r)} \neg c. \quad (3)$$

4. Otherwise, the set $\text{Def}_P(a)$ contains at least two rules (1) with $n + m > 0$ and

$$a \leftrightarrow \bigvee_{r \in \text{Def}_P(a)} \text{bd}_r \quad (4)$$

is introduced using a new atom bd_r for each $r \in \text{Def}_P(a)$ together with a formula

$$\text{bd}_r \leftrightarrow \bigwedge_{b \in \text{bd}^+(r)} b \wedge \bigwedge_{c \in \text{bd}^-(r)} \neg c. \quad (5)$$

The rest of the translation exploits the SCCs of the positive dependency graph of P that was defined in Section 2.1. The motivation is to limit the scope of ranking constraints which favors the length of the resulting translation. In particular, singleton components $\text{SCC}(a) = \{a\}$ require no special treatment if *tautological* rules with $a \in \{b_1, \dots, b_n\}$ in (1) have been removed. Plain completion (2) is sufficient for atoms involved in such components. However, for each atom $a \in \text{At}(P)$ having a non-trivial component $\text{SCC}(a)$ in $\text{DG}^+(P)$ such that $|\text{SCC}(a)| > 1$, two new atoms ext_a and int_a are introduced to formalize the *external* and *internal* support for a , respectively. These atoms are defined in terms of equivalences

$$\text{ext}_a \leftrightarrow \bigvee_{r \in \text{Ext}_P(a)} \text{bd}_r \quad (6)$$

$$\text{int}_a \leftrightarrow \bigvee_{r \in \text{Int}_P(a)} [\text{bd}_r \wedge \bigwedge_{b \in \text{bd}^+(r) \cap \text{SCC}(a)} (x_b <_m x_a)] \quad (7)$$

where x_a and x_b are bit vectors of width $m = \lceil \log_2(|\text{SCC}(a)| + 1) \rceil$ introduced for all atoms involved in $\text{SCC}(a)$. The formulas (6) and (7) are called *weak ranking constraints* and they are accompanied by

$$a \rightarrow \text{ext}_a \vee \text{int}_a, \quad (8)$$

$$\neg \text{ext}_a \vee \neg \text{int}_a. \quad (9)$$

Moreover, when $\text{Ext}_P(a) \neq \emptyset$ and the atom a happens to gain external support from these rules, the value of x_a is fixed to 0 by including the formula

$$\text{ext}_a \rightarrow (x_a =_m \bar{0}). \quad (10)$$

Example 5. Recall the program P from Example 1. The completion $\text{CC}(P)$ is:

$$\begin{aligned} a &\leftrightarrow \text{bd}_1 \vee \text{bd}_2. & \text{bd}_1 &\leftrightarrow b \wedge c. & \text{bd}_2 &\leftrightarrow d. \\ b &\leftrightarrow \text{bd}_3 \vee \text{bd}_4. & \text{bd}_3 &\leftrightarrow a \wedge \neg d. & \text{bd}_4 &\leftrightarrow a \wedge \neg c. \\ c &\leftrightarrow \neg d. \\ d &\leftrightarrow \neg c. \end{aligned}$$

Since P has only one non-trivial SCC, i.e., the component $\text{SCC}(a) = \text{SCC}(b) = \{a, b\}$, the weak ranking constraints resulting in $\text{R}(P)$ are

$$\begin{aligned} \text{ext}_a &\leftrightarrow \text{bd}_2. & \text{int}_a &\leftrightarrow \text{bd}_1 \wedge (x_b <_2 x_a). \\ \text{ext}_b &\leftrightarrow \perp. \\ \text{int}_b &\leftrightarrow [\text{bd}_3 \wedge (x_a <_2 x_b)] \vee [\text{bd}_4 \wedge (x_a <_2 x_b)]. \end{aligned}$$

In addition to these, the formulas

$$\begin{aligned} a &\rightarrow \text{ext}_a \vee \text{int}_a. & \neg \text{ext}_a \vee \neg \text{int}_a. & \text{ext}_a \rightarrow (x_a =_2 \bar{0}). \\ b &\rightarrow \text{ext}_b \vee \text{int}_b. & \neg \text{ext}_b \vee \neg \text{int}_b. \end{aligned}$$

are also included in $\text{R}(P)$. ■

Weak ranking constraints are sufficient whenever the goal is to compute only one answer set, or to check the existence of answer sets. However, they do not guarantee a one-to-one correspondence between the elements of $\text{AS}(P)$ and the set of models obtained for the translation $\text{BV}(P)$. To address this discrepancy, and to potentially make the computation of all answer sets or counting the number of answer sets more effective, *strong* ranking constraints can be imported from [20] as well. Actually, there are two mutually compatible variants of strong ranking constraints:

$$\text{bd}_r \rightarrow \bigvee_{b \in \text{bd}^+(r) \cap \text{SCC}(a)} \neg(x_b +_m \bar{1} <_m x_a) \quad (11)$$

$$\text{int}_a \rightarrow \bigvee_{r \in \text{Int}_P(a)} [\text{bd}_r \wedge \bigvee_{b \in \text{bd}^+(r) \cap \text{SCC}(a)} (x_a =_m x_b +_m \bar{1})]. \quad (12)$$

The *local* strong ranking constraint (11) is introduced for each $r \in \text{Int}_P(a)$. It is worth pointing out that the condition $\neg(x_b +_m \bar{1} <_m x_a)$ is equivalent to $x_b +_m \bar{1} \geq_m x_a$.⁶ On the other hand, the *global* variant (12) covers the internal support of a entirely. Finally, in order to prune copies of models of the translation that would correspond to the exactly same answer set of the original program, a formula

$$\neg a \rightarrow (x_a =_m \bar{0}) \quad (13)$$

is included for every atom a involved in a non-trivial SCC. We write $\text{R}^l(P)$ and $\text{R}^g(P)$ for the respective extensions of $\text{R}(P)$ with local/global strong ranking constraints, and $\text{R}^{\text{lg}}(P)$ obtained using both. Similar conventions are applied to $\text{BV}(P)$ to distinguish four variants in total. The correctness of these translations is addressed next.

Theorem 1. *Let P be a normal program and $\text{BV}(P)$ its bit-vector translation.*

1. *If S is an answer set of P , then there is a model $\langle M, \tau \rangle$ of $\text{BV}(P)$ such that $S = M \cap \text{At}(P)$.*
2. *If $\langle M, \tau \rangle$ is a model of $\text{BV}(P)$, then $S = M \cap \text{At}(P)$ is an answer set of P .*

Proof. To establish the correspondence of answer sets and models as formalized above, we appeal to the analogous property of the translation of P into difference logic (DL), denoted here by $\text{DL}(P)$. In DL, theory atoms $x \leq y + k$ constrain the difference of two integer variables x and y . Models can be represented as pairs $\langle I, \tau \rangle$ where I is a propositional interpretation and τ maps constants of theory atoms to integers so that $\langle I, \tau \rangle \models x \leq y + k \iff \tau(x) \leq \tau(y) + k$. The rest is analogous to Definition 2.

⁶ However, the form in (11) is used in our implementation, since $+_m$ and $<_m$ are amongst the base operators of the BOOLECTOR system.

(\implies) Suppose that S is an answer set of P . Then the results of [20] imply that there is a model $\langle M, \tau \rangle$ of $DL(P)$ such that $S = M \cap \text{At}(P)$. The valuation τ is condensed for each non-trivial SCC S of $DG^+(P)$ as follows. Let us partition S into $S_0 \sqcup \dots \sqcup S_n$ such that (i) $\tau(x_a) = \tau(x_b)$ for each $0 \leq i \leq n$ and $a, b \in S_i$, (ii) $\tau(x_a) = \tau(z)^7$ for each $a \in S_0$, and (iii) for each $0 \leq i < j \leq n$, $a \in S_i$, and $b \in S_j$, $\tau(x_a) \leq \tau(x_b)$. Then define τ' for the bit vector x_a associated with an atom $a \in S_i$ by setting $\tau'(x_a, j) = 1$ iff the j^{th} bit of \bar{i} is 1, i.e., $\tau'(x_a) = \bar{i}$. It follows that $\langle I, \tau \rangle \models x_b \leq x_a - 1$ iff $\langle I, \tau' \rangle \models x_b <_m x_a$ for any $a, b \in S$. Moreover, we have $\langle M, \tau \rangle \models (x_a \leq z + 0) \wedge (z \leq x_a + 0)$ iff $\langle M, \tau' \rangle \models x_a =_m \bar{0}$ for any $a \in S$. Due to the similar structures of $DL(P)$ and $BV(P)$, we obtain $\langle M, \tau \rangle \models BV(P)$ as desired.

(\Leftarrow) Let $\langle M, \tau \rangle$ be a model of $BV(P)$. Then define τ' such that $\tau'(x) = \sum_{i=1}^{w(x)} (\tau(x, i) \cdot 2^{w(x)-i})$ where x on the left hand side stands for the integer variable corresponding to the bit vector x on the right hand side. It follows that $\langle I, \tau \rangle \models x_b <_m x_a$ iff $\langle I, \tau' \rangle \models x_b \leq x_a - 1$. By setting $\tau'(z) = 0$, we obtain $\langle M, \tau \rangle \models x_a =_m \bar{0}$ if and only if $\langle M, \tau' \rangle \models (x_a \leq z + 0) \wedge (z \leq x_a + 0)$. The strong analogy present in the structures of $BV(P)$ and $DL(P)$ implies that $\langle M, \tau' \rangle$ is a model of $DL(P)$. Thus, $S = M \cap \text{At}(P)$ is an answer set of P by [20]. \square

Even tighter relationships of answer sets and models can be established for the translations $BV^1(P)$, $BV^g(P)$, and $BV^{lg}(P)$. It can be shown that the model $\langle M, \tau \rangle$ of $BV^*(P)$ corresponding to an answer set S of P is unique, i.e., there is no other model $\langle N, \tau' \rangle$ of the translation such that $S = N \cap \text{At}(P)$. These results contrast with [20]: the analogous extensions $DL^*(P)$ guarantee the uniqueness of M in a model $\langle M, \tau \rangle$ but there are always infinitely many copies $\langle M, \tau' \rangle$ of $\langle M, \tau \rangle$ such that $\langle M, \tau' \rangle \models DL^*(P)$. Such a valuation τ' can be simply obtained by setting $\tau'(x) = \tau(x) + 1$ for any x .

4 Native Support for Extended Rule Types

The input syntax of the SMOBELS system was soon extended by further rule types [28]. In solver interfaces, the rule types usually take the following simple syntactic forms:

$$\{a_1, \dots, a_l\} \leftarrow b_1, \dots, b_n, \sim c_1, \dots, \sim c_m. \quad (14)$$

$$a \leftarrow l\{b_1, \dots, b_n, \sim c_1, \dots, \sim c_m\}. \quad (15)$$

$$a \leftarrow l\{b_1 = w_{b_1}, \dots, b_n = w_{b_n}, \sim c_1 = w_{c_1}, \dots, \sim c_m = w_{c_m} = w_{c_m}\}. \quad (16)$$

The body of a *choice rule* (14) is interpreted in the same way as that of a normal rule (1). The head, in contrast, allows to derive any subset of atoms a_1, \dots, a_l , if the body is satisfied, and to make a *choice* in this way. The head a of a *cardinality rule* (15) is derived, if its body is satisfied, i.e., the number of satisfied literals amongst b_1, \dots, b_n and $\sim c_1, \dots, \sim c_m$ is at least l acting as the *lower bound*. A *weight rule* of the form (16) generalizes this idea by assigning arbitrary positive weights to literals (rather than 1s). The body is satisfied if the sum of weights assigned to satisfied literals is at least l , thus enabling one to infer the head a using the rule. In practise, the grounding components used in ASP systems allow for more versatile use of cardinality and weight rules, but the primitive forms (14), (15), and (16) provide a solid basis for efficient implementation via translations. The reader is referred to [28] for a generalization of answer sets for programs involving such extended rule types. The respective class of *weight constraint programs* (WCPs) is typically supported by SMOBELS compatible systems.

Whenever appropriate, it is possible to translate extended rule types as introduced above back to normal rules. To this end, a number of transformations are addressed in [19] and they have been implemented as a tool called LP2NORMAL⁸. For instance, the head of a choice rule (14) can be captured in terms of rules

$$\begin{aligned} a_1 &\leftarrow b, \sim \bar{a}_1. & \dots & a_l \leftarrow b, \sim \bar{a}_l. \\ \bar{a}_1 &\leftarrow \sim a_1. & \dots & \bar{a}_l \leftarrow \sim a_l. \end{aligned}$$

where $\bar{a}_1, \dots, \bar{a}_l$ are new atoms and b is a new atom standing for the body of (14) which can be defined using (14) with the head replaced by b . We assume that this transformation is applied at first to remove

⁷ A special variable z is used as a placeholder for the constant 0 in the translation $DL(P)$ [20].

⁸ <http://www.tcs.hut.fi/Software/asptools/>

```
gringo program.lp instance.lp \
| smodels -internal -nolookahead \
| lpcat -s=symbols.txt \
| lp2bv [-l] [-g] \
| boolector -fm
```

Fig. 1. Unix shell pipeline for running a benchmark instance

choice rules when the goal is to translate extended rule types into bit-vector logic. The strength of this transformation is locality, i.e., it can be applied on a rule-by-rule basis, and linearity with respect to the length of the original rule (14). To the contrary, linear normalization of cardinality and weight rules seems impossible. Thus, we also provide direct translations into formulas of bit-vector logic.

We present the translation of a weight rule (16) whereas the translation of a cardinality rule (15) is obtained as a special case $w_{b_1} = \dots = w_{b_n} = w_{c_1} = \dots = w_{c_m} = 1$. The body of a weight rule can be evaluated using bit vectors s_1, \dots, s_{n+m} of width $k = \lceil \log_2(\sum_{i=1}^n w_{b_i} + \sum_{i=1}^m w_{c_i} + 1) \rceil$ constrained by $2 \times (n + m)$ formulas

$$\begin{array}{ll}
b_1 \rightarrow (s_1 =_k \overline{w_{b_1}}), & \neg b_1 \rightarrow (s_1 =_k \overline{0}), \\
b_2 \rightarrow (s_2 =_k s_1 +_k \overline{w_{b_2}}), & \neg b_2 \rightarrow (s_2 =_k s_1), \\
\vdots & \vdots \\
b_n \rightarrow (s_n =_k s_{n-1} +_k \overline{w_{b_n}}), & \neg b_n \rightarrow (s_n =_k s_{n-1}), \\
c_1 \rightarrow (s_{n+1} =_k s_n), & \neg c_1 \rightarrow (s_{n+1} =_k s_n +_k \overline{w_{c_1}}), \\
\vdots & \vdots \\
c_m \rightarrow (s_{n+m} =_k s_{n+m-1}), & \neg c_m \rightarrow (s_{n+m} =_k s_{n+m-1} +_k \overline{w_{c_m}}).
\end{array}$$

The lower bound l of (16) can be checked in terms of the formula $\neg(s_{n+m} <_k \bar{l})$ where we assume that \bar{l} is of width k , since the rule can be safely deleted otherwise. In view of the overall translation, the formula $\text{bd}_r \leftrightarrow \neg(s_{n+m} <_k \bar{l})$ can be used in conjunction with the completion formula (4). Weight rules also contribute to the dependency graph $\text{DG}^+(P)$ in analogy to normal rules, i.e., the head a depends on all positive body atoms b_1, \dots, b_n . In this way, $\text{BV}(P)$ generalizes for programs P having extended rules.

5 Experimental Results

A new translator called LP2BV was implemented as a derivative of LP2DIFF⁹ that translates logic programs into difference logic. In contrast, the new translator will provide its output in the bit-vector format. In analogy to its predecessor, it expects to receive its input in the SMOLELS¹⁰ file format. Models of the resulting bit-vector theory are searched for using BOOLECTOR¹¹ (v. 1.4.1) [6] and Z3¹² (v. 2.11) [9] as back-end solvers. The goal of our preliminary experiments was to see how the performances of systems based on LP2BV compare with the performance of a state-of-the-art ASP solver CLASP¹³ (v. 1.3.5) [13]. The experiments were based on the NP-complete benchmarks of the ASP Competition 2009. In this benchmark collection, there are 23 benchmark problems with 516 instances in total. Before invoking a translator and the respective SMT solver, we performed a few preprocessing steps, as detailed in Figure 1, by calling:

- GRINGO (v. 2.0.5), for grounding the problem encoding and a given instance;
- SMOLELS¹⁴ (v. 2.34), for simplifying the resulting ground program;

⁹ <http://www.tcs.hut.fi/Software/lp2diff/>

¹⁰ <http://www.tcs.hut.fi/Software/smodels/>

¹¹ <http://fmv.jku.at/boolector/>

¹² <http://research.microsoft.com/en-us/um/redmond/projects/z3/>

¹³ <http://www.cs.uni-potsdam.de/clasp/>

¹⁴ <http://www.tcs.hut.fi/Software/smodels/>

Table 1. Experimental results without normalization

Benchmark	INST	CLASP	LP2BV+BOOLECTOR				LP2BV+Z3				LP2DIFF+Z3			
			W	L	G	LG	W	L	G	LG	W	L	G	LG
Overall Performance	516	465 347/118	276 188/ 88	244 161/ 83	261 174/ 87	256 176/ 80	217 142/ 75	216 147/ 69	194 124/ 70	204 135/ 69	360 257/103	349 251/ 98	324 225/ 99	324 226/ 98
KnightTour	10	8/ 0	2/ 0	1/ 0	0/ 0	0/ 0	1/ 0	0/ 0	0/ 0	1/ 0	6/ 0	6/ 0	4/ 0	5/ 0
GraphColouring	29	8/ 0	7/ 0	7/ 0	7/ 0	7/ 0	6/ 0	7/ 0	7/ 0	7/ 0	7/ 0	7/ 0	7/ 0	7/ 0
WireRouting	23	11/11	2/ 3	1/ 1	1/ 2	0/ 2	1/ 3	0/ 0	0/ 0	0/ 1	3/ 3	2/ 3	2/ 4	5/ 3
DisjunctiveScheduling	10	5/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0
GraphPartitioning	13	6/ 7	3/ 0	3/ 0	3/ 0	3/ 0	4/ 0	4/ 0	4/ 0	3/ 0	6/ 2	6/ 1	6/ 1	6/ 1
ChannelRouting	11	6/ 2	6/ 2	6/ 2	6/ 2	6/ 2	5/ 2	6/ 2	6/ 2	6/ 2	6/ 2	6/ 2	6/ 2	6/ 2
Solitaire	27	19/ 0	2/ 0	5/ 0	1/ 0	4/ 0	0/ 0	0/ 0	0/ 0	0/ 0	21/ 0	21/ 0	20/ 0	21/ 0
Labyrinth	29	26/ 0	1/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0
WeightBoundedDominatingSet	29	26/ 0	18/ 0	18/ 0	17/ 0	18/ 0	12/ 0	12/ 0	11/ 0	12/ 0	22/ 0	22/ 0	22/ 0	21/ 0
MazeGeneration	29	10/15	8/15	1/15	0/15	0/16	5/16	1/15	0/15	1/15	10/17	10/15	5/15	4/15
15Puzzle	16	16/ 0	16/ 0	15/ 0	14/ 0	15/ 0	4/ 0	4/ 0	5/ 0	5/ 0	0/ 0	0/ 0	0/ 0	0/ 0
BlockedNQueens	29	15/14	2/ 2	0/ 2	1/ 2	0/ 2	1/ 0	2/ 0	2/ 0	0/ 0	15/13	15/13	15/12	15/13
ConnectedDominatingSet	21	10/10	10/11	9/ 8	10/11	6/ 3	10/10	9/10	10/ 9	10/ 9	9/ 8	7/ 6	9/ 7	7/ 6
EdgeMatching	29	29/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	0/ 0	3/ 0	1/ 0	3/ 0	2/ 0
Fastfood	29	10/19	9/16	10/16	10/16	9/16	9/ 9	9/ 9	9/10	9/ 9	10/18	10/18	10/18	10/18
GeneralizedSlitherlink	29	29/ 0	29/ 0	20/ 0	29/ 0	29/ 0	29/ 0	29/ 0	16/ 0	29/ 0	29/ 0	29/ 0	29/ 0	29/ 0
HamiltonianPath	29	29/ 0	27/ 0	25/ 0	29/ 0	28/ 0	26/ 0	27/ 0	25/ 0	26/ 0	29/ 0	29/ 0	29/ 0	29/ 0
Hanoi	15	15/ 0	15/ 0	15/ 0	15/ 0	15/ 0	5/ 0	5/ 0	5/ 0	4/ 0	15/ 0	15/ 0	15/ 0	15/ 0
HierarchicalClustering	12	8/ 4	8/ 4	8/ 4	8/ 4	8/ 4	4/ 4	4/ 4	4/ 4	4/ 4	8/ 4	8/ 4	8/ 4	8/ 4
SchurNumbers	29	13/16	6/16	5/16	5/16	5/16	9/16	9/16	9/16	9/16	11/16	11/16	11/16	11/16
Sokoban	29	9/20	9/19	8/19	8/19	8/19	7/15	7/13	7/14	5/13	9/20	9/20	9/20	9/20
Sudoku	10	10/ 0	5/ 0	4/ 0	4/ 0	5/ 0	4/ 0	4/ 0	4/ 0	4/ 0	9/ 0	8/ 0	8/ 0	9/ 0
TravellingSalesperson	29	29/ 0	3/ 0	0/ 0	6/ 0	10/ 0	0/ 0	8/ 0	0/ 0	0/ 0	29/ 0	29/ 0	7/ 0	7/ 0

- LPCAT (v. 1.18), for removing all unused atom numbers, for making the atom table of the ground program contiguous, and for extracting the symbols for later use; and
- LP2NORMAL (version 1.11), for normalizing the program.

The last step is optional and not included as part of the pipeline in Figure 1. Pipelines of this kind were executed under Linux/Ubuntu operating system running on six-core AMD Opteron^(TM) 2435 processors under 2.6 GHz clock rate and with 2.7 GB memory limit that corresponds to the amount of memory available in the ASP Competition 2009.

For each system based on a translator and a back-end solver, there are four variants of the system to consider: W indicates that only weak ranking constraints are used, while L, G, and LG mean that either local, or global, or both local and global strong ranking constraints, respectively, are employed when translating the logic program.

Table 1 collects the results from our experiments without normalization whereas Table 2 shows the results when LP2NORMAL [19] was used to remove extended rule types discussed in Section 4. In both tables, the first column gives the name of the benchmark, followed by the number of instances of that particular benchmark in the second column. The following columns indicate the numbers of instances that were solved by the systems considered in our experiments. A notation like 8/4 means that the system was able to solve eight satisfiable and four unsatisfiable instances in that particular benchmark. Hence, if there are 15 instances in a benchmark and the system could only solve 8/4, this means that the system was unable to solve the remaining three instances within the time limit of 600 seconds, i.e. ten minutes, per instance¹⁵. As regards the number of solved instances in each benchmark, the best performing translation-based approaches are highlighted in boldface. Though it was not shown in all tables, we also run the experiments using translator LP2DIFF with Z3 as back-end solver, and the summary is included in Table 3—giving an overview of experimental results in terms of total numbers of instances solved out of 516.

It is apparent that the systems based on LP2BV did not perform very well without normalization. As indicated by Table 3, the overall performance was even worse than that of systems using LP2DIFF for

¹⁵ One observation is that the performance of systems based on LP2BV is quite stable: even when we extended the time limit to 20 minutes, the results did not change much (differences of only one or two instances were perceived in most cases).

Table 2. Experimental results with normalization

Benchmark	INST	CLASP	LP2BV+BOOLECTOR				LP2BV+Z3			
			W	L	G	LG	W	L	G	LG
Overall Performance	516	459 346/113	381 279/102	343 243/100	379 278/101	381 281/100	346 240/106	330 231/ 99	325 224/101	331 232/ 99
KnightTour	10	10/ 0	2/0	2/0	1/ 0	0/ 0	1/ 0	0/ 0	0/ 0	0/ 0
GraphColouring	29	9/ 0	8/ 0	8/ 0	8/ 0	8/ 0	9/2	9/2	9/2	9/2
WireRouting	23	11/11	2/ 6	1/ 3	1/ 3	1/ 3	2/7	1/ 4	1/ 4	1/ 3
DisjunctiveScheduling	10	5/ 0	5/0	5/0	5/0	5/0	5/0	5/0	5/0	5/0
GraphPartitioning	13	4/ 1	5/0	5/0	4/ 0	5/0	2/ 1	2/ 1	2/ 1	2/ 0
ChannelRouting	11	6/ 2	6/2	6/2	6/2	6/2	6/2	6/2	6/2	6/2
Solitaire	27	18/ 0	23/0	23/0	23/0	23/0	22/ 0	22/ 0	22/ 0	22/ 0
Labyrinth	29	27/ 0	1/ 0	1/ 0	2/ 0	3/0	0/ 0	0/ 0	0/ 0	0/ 0
WeightBoundedDominatingSet	29	25/ 0	15/ 0	15/ 0	15/ 0	16/0	10/ 0	10/ 0	10/ 0	10/ 0
MazeGeneration	29	10/15	8/15	0/15	0/15	0/16	5/16	0/15	0/15	0/15
15Puzzle	16	15/ 0	16/0	16/0	16/0	16/0	11/ 0	10/ 0	11/ 0	11/ 0
BlockedNQueens	29	15/14	14/14	14/14	14/14	14/14	15/14	15/14	15/14	15/14
ConnectedDominatingSet	21	10/11	10/11	8/11	9/11	9/10	10/11	9/11	9/11	9/11
EdgeMatching	29	29/ 0	29/ 0	29/ 0	29/ 0	29/ 0	29/ 0	29/ 0	29/ 0	29/ 0
Fastfood	29	10/19	9/14	9/15	9/16	9/15	0/13	0/10	0/12	0/12
GeneralizedSlitherlink	29	29/ 0	29/0	21/ 0	29/0	29/0	29/0	29/0	21/ 0	29/0
HamiltonianPath	29	29/ 0	29/0	28/ 0	29/0	29/0	29/0	29/0	29/0	29/0
Hanoi	15	15/ 0	15/0	15/0	15/0	15/0	15/0	15/0	15/0	15/0
HierarchicalClustering	12	8/ 4	8/4	8/4	8/4	8/4	8/4	8/4	8/4	8/4
SchurNumbers	29	13/16	10/16	10/16	9/16	10/16	13/16	13/16	13/16	13/16
Sokoban	29	9/20	9/20	9/20	9/20	9/20	9/20	9/20	9/20	9/20
Sudoku	10	10/ 0	10/0	10/0	10/0	10/0	10/0	10/0	10/0	10/0
TravellingSalesperson	29	29/ 0	16/ 0	0/ 0	27/0	27/0	0/ 0	0/ 0	0/ 0	0/ 0

translation and Z3 for model search. However, if the input was first translated into a normal logic program using LP2NORMAL, i.e., before translation into a bit-vector theory, the performance was clearly better. Actually, it exceeded that of the systems based on LP2DIFF and became closer to that of CLASP. We note that normalization does not help so much in case of LP2DIFF and the experimental results obtained using both normalized and unnormalized instances are quite similar in terms of solved instances. Thus it seems that solvers for bit-vector logic are not able to make the best of native translations of cardinality and weight rules from Section 4 in full. If an analogous translation into difference logic is used, as implemented in LP2DIFF, such a negative effect was not perceived using Z3. Our understanding is that the efficient graph-theoretic satisfiability check for difference constraints used in the search procedure of Z3 turns the native translation feasible as well. As indicated by our test results, BOOLECTOR is clearly better back-end solver for LP2BV than Z3. This was to be expected since BOOLECTOR is a native solver for bit-vector logic whereas Z3 supports a wider variety of SMT fragments and can be used for more general purposes. Moreover, the design of LP2BV takes into account operators of bit-vector logic which are directly supported by BOOLECTOR and not implemented as syntactic sugar.

In addition, we note on the basis of our results that the performance of the state-of-the-art ASP solver CLASP is significantly better, and the translation-based approaches to computing stable models are still left behind. By the results of Table 2, even the best variants of systems based on LP2BV did not work well enough to compete with CLASP. The difference is especially due to the following benchmarks: *Knight Tour*, *Wire Routing*, *Graph Partitioning*, *Labyrinth*, *Weight Bounded Dominating Set*, *Fastfood*, and *Travelling Salesperson*. All of them involve either recursive rules (*Knight Tour*, *Wire Routing*, and *Labyrinth*), weight rules (*Weight Bounded Dominating Set* and *Fastfood*), or both (*Graph Partitioning* and *Travelling Salesperson*). Hence, it seems that handling recursive rules and weight constraints in the translational approach is less efficient compared to their native implementation in CLASP. When using the current normalization techniques to remove cardinality and weight rules, the sizes of ground programs tend to increase significantly and, in particular, if weight rules are abundant. For example, after normalization the ground programs are ten times larger for the benchmark *Weight Bounded Dominating Set*, and five times larger for *Fastfood*. It is also worth pointing out that the efficiency of CLASP turned out to be insensitive to normalization.

Table 3. Summary of the experimental results

System	W	L	G	LG
LP2BV+BOOLECTOR	276	244	261	256
LP2BV+Z3	217	216	194	204
LP2DIFF+Z3	360	349	324	324
CLASP	465			
LP2NORMAL2BV+BOOLECTOR	381	343	379	381
LP2NORMAL2BV+Z3	346	330	325	331
LP2NORMAL2DIFF+Z3	364	357	349	349
LP2NORMAL+CLASP	459			

While having trouble with recursive rules and weight constraints for particular benchmarks, the translational approach handles certain large instances quite well. The largest instances in the experiments belong to the *Disjunctive Scheduling* benchmark, of which all instances are ground programs of size over one megabyte but after normalization¹⁶, the LP2BV systems can solve as many instances as CLASP.

6 Conclusion

In this paper, we present a novel and concise translation from normal logic programs into fixed-width bit-vector theories. Moreover, the extended rule types supported by SMOLELS compatible answer set solvers can be covered via native translations. The length of the resulting translation is linear with respect to the length of the original program. The translation has been implemented as a translator, LP2BV, which enables the use of bit-vector solvers in the search for answer sets. Our preliminary experimental results indicate a level of performance which is similar to that obtained using solvers for difference logic. However, this presumes one first to translate extended rule types into normal rules and then to apply the translation into bit-vector logic. One potential explanation for such behavior is the way in which SMT solvers implement reasoning with bit vectors: a predominant strategy is to translate theory atoms involving bit vectors into propositional formulas and to apply satisfiability checking techniques systematically. We anticipate that an improved performance could be obtained if a native support for certain bit vector primitives were incorporated into SMT solvers directly. When comparing to the state-of-the-art ASP solver CLASP, we noticed that the performance of the translation based approach compared unfavorably, in particular, for benchmarks which contained recursive rules or weight constraints or both. This indicates that the performance can be improved by developing new translation techniques for these two features. In order to obtain a more comprehensive view of the performance characteristics of the translational approach, the plan is to extend our experimental setup to include benchmarks that were used in the third ASP competition [7]. Moreover, we intend to use the new SMT library format [4] in future versions of our translators.

Acknowledgments This research has been partially funded by the Academy of Finland under the project “*Methods for Constructing and Solving Large Constraint Models*” (MCM, #122399).

References

1. Krzysztof Apt, Howard Blair, and Adrian Walker. Towards a theory of declarative knowledge. In *Foundations of Deductive Databases and Logic Programming*, pages 89–148. Morgan Kaufmann, 1988.
2. Marcello Balduccini. Industrial-size scheduling with ASP+CP. In Delgrande and Faber [10], pages 284–296.
3. Clark Barrett, Roberto Sebastiani, Sanjit Seshia, and Cesare Tinelli. Satisfiability modulo theories. In Biere et al. [5], pages 825–885.
4. Clark Barrett, Aaron Stump, and Cesare Tinelli. The SMT-LIB standard version 2.0.

¹⁶ In this benchmark, normalization does not affect the size of grounded programs significantly.

5. Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors. *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2009.
6. Robert Brummayer and Armin Biere. Boolector: An efficient SMT solver for bitvectors and arrays. In Stefan Kowalewski and Anna Philippou, editors, *TACAS*, volume 5505 of *Lecture Notes in Computer Science*, pages 174–177. Springer, 2009.
7. Francesco Calimeri, Giovambattista Ianni, Francesco Ricca, Mario Alviano, Annamaria Bria, Gelsomina Catalano, Susanna Cozza, Wolfgang Faber, Onofrio Febbraro, Nicola Leone, Marco Manna, Alessandra Martello, Claudio Panetta, Simona Perri, Kristian Reale, Maria Carmela Santoro, Marco Sirianni, Giorgio Terracina, and Pierfrancesco Veltri. The third answer set programming competition: Preliminary report of the system competition track. In Delgrande and Faber [10], pages 388–403.
8. Keith Clark. Negation as failure. In *Logic and Data Bases*, pages 293–322, 1977.
9. Leonardo Mendonça de Moura and Nikolaj Bjørner. Z3: An efficient SMT solver. In *TACAS*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer, 2008.
10. James Delgrande and Wolfgang Faber, editors. *Logic Programming and Nonmonotonic Reasoning - 11th International Conference, LPNMR 2011, Vancouver, Canada, May 16-19, 2011. Proceedings*, volume 6645 of *Lecture Notes in Computer Science*. Springer, 2011.
11. Marc Denecker, Joost Vennekens, Stephen Bond, Martin Gebser, and Mirosław Truszczyński. The second answer set programming competition. In Erdem et al. [12], pages 637–654.
12. Esra Erdem, Fangzhen Lin, and Torsten Schaub, editors. *Logic Programming and Nonmonotonic Reasoning, 10th International Conference, LPNMR 2009, Potsdam, Germany, September 14-18, 2009. Proceedings*, volume 5753 of *Lecture Notes in Computer Science*. Springer, 2009.
13. Martin Gebser, Benjamin Kaufmann, André Neumann, and Torsten Schaub. *clasp*: A conflict-driven answer set solver. In Chitta Baral, Gerhard Brewka, and John Schlipf, editors, *LPNMR*, volume 4483 of *Lecture Notes in Computer Science*, pages 260–265. Springer, 2007.
14. Martin Gebser, Max Ostrowski, and Torsten Schaub. Constraint answer set solving. In Patricia Hill and David Scott Warren, editors, *ICLP*, volume 5649 of *Lecture Notes in Computer Science*, pages 235–249. Springer, 2009.
15. Michael Gelfond and Nicola Leone. Logic programming and knowledge representation – the A-Prolog perspective. *Artif. Intell.*, 138(1-2):3–38, 2002.
16. Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programming. In *ICLP/SLP*, pages 1070–1080, 1988.
17. Tomi Janhunen. Some (in)translatability results for normal logic programs and propositional theories. *Journal of Applied Non-Classical Logics*, 16(1–2):35–86, June 2006.
18. Tomi Janhunen, Guohua Liu, and Ilkka Niemelä. Tight integration of non-ground answer set programming and satisfiability modulo theories. In Eugenia Ternovska and David Mitchell, editors, *Working Notes of Grounding and Transformations for Theories with Variables*, pages 1–13, Vancouver, Canada, May 2011.
19. Tomi Janhunen and Ilkka Niemelä. Compact translations of non-disjunctive answer set programs to propositional clauses. In Marcello Balduccini and Tran Cao Son, editors, *Logic Programming, Knowledge Representation, and Nonmonotonic Reasoning*, volume 6565 of *Lecture Notes in Computer Science*, pages 111–130. Springer, 2011.
20. Tomi Janhunen, Ilkka Niemelä, and Mark Sevalnev. Computing stable models via reductions to difference logic. In Erdem et al. [12], pages 142–154.
21. Victor Marek and Venkatramana Subrahmanian. The relationship between stable, supported, default and autoepistemic semantics for general logic programs. *Theor. Comput. Sci.*, 103(2):365–386, 1992.
22. Victor Marek and Mirek Truszczyński. Stable models and an alternative logicprogramming paradigm. In *The Logic Programming Paradigm: A 25-Year Perspective*, pages 375–398. Springer, 1999.
23. Veena Mellarkod and Michael Gelfond. Integrating answer set reasoning with constraint solving techniques. In Jacques Garrigue and Manuel Hermenegildo, editors, *FLOPS*, volume 4989 of *Lecture Notes in Computer Science*, pages 15–31. Springer, 2008.
24. Ilkka Niemelä. Logic programs with stable model semantics as a constraint programming paradigm. *Ann. Math. Artif. Intell.*, 25(3-4):241–273, 1999.
25. Ilkka Niemelä. Stable models and difference logic. *Ann. Math. Artif. Intell.*, 53(1-4):313–329, 2008.
26. Robert Nieuwenhuis and Albert Oliveras. DPLL(T) with exhaustive theory propagation and its application to difference logic. In Kousha Etessami and Sriram Rajamani, editors, *CAV*, volume 3576 of *Lecture Notes in Computer Science*, pages 321–334. Springer, 2005.
27. Silvio Ranise and Cesare Tinelli. The SMT-LIB format: An initial proposal.
28. Patrik Simons, Ilkka Niemelä, and Timo Soininen. Extending and implementing the stable model semantics. *Artif. Intell.*, 138(1-2):181–234, 2002.